Online Balanced Repartitioning*

Chen Avin[1] Andreas Loukas[2] Maciej Pacut[3] Stefan Schmid[4,2]

[1] Ben Gurion University of the Negev, Israel [2] TU Berlin, Germany
[3] University of Wroclaw, Poland [4] Aalborg University, Denmark

arXiv:1511.02074v3 [cs.DS] 18 Jul 2016

Abstract. This paper initiates the study of deterministic algorithms for collocating frequently communicating nodes in a distributed networked systems in an online fashion. In particular, we introduce the *Balanced RePartitioning* (BRP) problem: Given an arbitrary sequence of pairwise communication requests between n nodes, with patterns that may change over time, the objective is to dynamically partition the nodes into ℓ clusters, each of size k, at a minimum cost. Every communication request needs to be served: if the communicating nodes are located in the same cluster, the request is served locally, at cost 0; if the nodes are located in different clusters, the request is served remotely using inter-cluster communication, at cost 1. The partitioning can be updated dynamically (i.e., repartitioned), by *migrating* nodes between clusters at cost α per node migration. The goal is to devise online algorithms which find a good trade-off between the communication and the migration cost, i.e., "rent" or "buy", while maintaining partitions which minimize the number of inter-cluster communications. BRP features interesting connections to other well-known online problems. In particular, we show that scenarios with $\ell = 2$ generalize online paging, and scenarios with $k = 2$ constitute a novel online version of maximum matching. We consider settings both with and without cluster-size augmentation. Somewhat surprisingly, we prove that any deterministic online algorithm has a competitive ratio of at least k, *even with augmentation*. Our main technical contribution is an $O(k \log k)$-competitive deterministic algorithm for the setting with (constant) augmentation. This is attractive as, in contrast to ℓ, k is likely to be small in practice. For the case of matching ($k = 2$), we present a constant competitive algorithm that does not rely on augmentation.

Keywords: Dynamic Graphs, Clustering, Graph Partitioning, Algorithms, Competitive Analysis, Cloud Computing

1 Introduction

Graph partitioning problems, like minimum graph bisection and maximum matching, are among the most fundamental problems in Theoretical Computer Science. Due to their numerous practical applications (e. g., communication networks, data mining, social networks, etc. [4,14,26]), partitioning problems are also among the most intensively studied problems. Interestingly however, not much is known today about how to dynamically partition nodes which interact or communicate in a time-varying fashion.

This paper initiates the study of a natural model for *online* graph partitioning. We are given a set of n nodes with time-varying pairwise communication patterns, which we have to partition into ℓ clusters of equal size k. Intuitively, we would like to minimize inter-cluster links by mapping frequently communicating nodes to the same cluster. Since communication patterns change over time, partitions should be dynamically readjusted, that is, the nodes should be *repartitioned*, in an online manner, by *migrating* them between clusters. The objective is to jointly minimize inter-cluster communication and reconfiguration costs, defined respectively as the number of communication requests served remotely and the number of times nodes are migrated from one cluster to another.

One practical motivation for our problem arises in the context of server virtualization in datacenters. Distributed cloud applications, including batch processing applications such as MapReduce, streaming applications such as Apache Flink or Apache Spark, and scale-out databases and key-value stores such as Cassandra, generate a significant amount of network traffic and a considerable fraction of their runtime is due to network activity [24]. For example, traces of jobs from a Facebook cluster reveal that network transfers on average account for 33% of the execution time [13]. In such applications, it is desirable that frequently communicating virtual machines are *collocated*, i.e., mapped to the same physical server, since communication across the network

* Technical report of the corresponding DISC 2016 paper. Research supported by the German-Israeli Foundation for Scientific Research (GIF) Grant I-1245-407.6/2014 and by the Polish National Science Centre grant DEC-2013/09/B/ST6/01538.

(i.e., inter-server communication) induces network load and latency. However, migrating virtual machines between servers also comes at a price: the state transfer is bandwidth intensive, and may even lead to short service interruptions. Therefore the goal is to design online algorithms which find a good trade-off between the inter-server communication cost and the migration cost, similar in spirit to classical ski rental and rent-or-buy problems.

Formally, the inputs to the *Balanced RePartitioning* (BRP) problem are:

1. A set V of $n = |V|$ nodes (e.g., the virtual machines), initially distributed arbitrarily across ℓ clusters $\mathcal{C} = \{C_1, \ldots, C_\ell\}$ (e.g., the physical servers, interconnected by a top-of-the-rack switch [2]), each of size k (e.g., the number of cores or slots for virtual machines).
2. An arbitrary and possibly infinite sequence σ of $|\sigma|$ communication requests, $\sigma = \{u_1, v_1\}, \{u_2, v_2\}, \{u_3, v_3\}, \ldots, \{u_{|\sigma|}, v_{|\sigma|}\}$. For any t, $\sigma_t = \{u_t, v_t\}$ denotes a communication request: at time t, nodes $u_t, v_t \in V$ exchange (a fixed amount of) data. Intuitively, every request σ_t can be thought of as an edge of the communication graph which appears at time t and then disappears at $t + 1$.

At any time t, each node $v \in V$ is assigned to a cluster, which we will refer to by $C_t(v) \in \mathcal{C}$. If the time t is clear from the context or irrelevant, we simply write $C(v)$. We call two nodes $u, v \in V$ *collocated* if they are in the same cluster: $C(u) = C(v)$. We consider two settings:

1. *Without augmentation:* The nodes fit perfectly into the clusters, i.e., $n = k \cdot \ell$.
2. *With augmentation:* The online algorithm has access to additional space in each cluster. We say that an algorithm is δ-augmented if the size of each cluster is $k' = \delta \cdot k$, whereas the total number of nodes remains $n = k \cdot \ell < k' \cdot \ell$. As usual, in the competitive analysis, the augmented online algorithm is compared to the optimal offline algorithm without augmentation.

At each time t, the online algorithm needs to serve the communication request $\{u_t, v_t\}$, but can also repartition the nodes into new clusters before serving the request. We assume that a communication request between two nodes located in different clusters costs 1, a communication request between two nodes collocated in the same cluster costs 0, and migrating a node from one cluster to another costs $\alpha \geq 1$. Note that in a setting without augmentation, due to cluster size constraints, a node can never be migrated alone, but it must be swapped with another node at a cost of (at least) 2α.

As it turns out, BRP features some interesting connections to other well-known graph and online problems: (i) The static version (without migration) is the minimum balanced graph partitioning problem (where ℓ is the number of components). (ii) For $\ell = 2$, BRP can be shown to be a generalization of online paging, where the first cluster simulates the cache (the small but fast memory) and the second the slow but large memory. (iii) For $k = 2$, BRP is a novel online version of maximum matching. In the static case, maximum matching is a special case of minimum balanced graph partitioning with $n/2$ components.

The cost of an algorithm ALG for a given sequence of communication requests σ is

$$\mathrm{ALG}(\sigma) = \sum_{t=1}^{|\sigma|} \mathrm{mig}(\sigma_t; \mathrm{ALG}) + \mathrm{com}(\sigma_t; \mathrm{ALG}), \tag{1}$$

where $\mathrm{mig}(\sigma_t; \mathrm{ALG})$ is the migration cost at time t (α or 0) and $\mathrm{com}(\sigma_t; \mathrm{ALG})$ is the communication cost of σ_t (1 or 0). Let $\textsc{On}(\sigma)$ and $\textsc{Off}(\sigma)$ be the cost induced by σ on an online algorithm On and an optimal offline algorithm Off, respectively. In contrast to On, which learns the requests one-by-one as it serves them, Off has a complete knowledge of the entire request sequence σ ahead of time. We are in the realm of online algorithms and competitive analysis: We want to design online repartitioning algorithms which provide conservative (worst-case) guarantees, and minimize the (strict) competitive ratio:

$$\rho(\textsc{On}) = \max_{\sigma} \frac{\textsc{On}(\sigma)}{\textsc{Off}(\sigma)}. \tag{2}$$

To be competitive, an online repartitioning algorithm has to define a strategy for each of the following questions:

A) *Serve remotely or migrate ("rent or buy")?* If a communication pattern is short-lived, it may not be worthwhile to collocate the nodes: the migration cost cannot be amortized.
B) *Where to migrate, and what?* If nodes should be collocated, the question becomes where. Should u_t be migrated to $C(v_t)$, v_t to $C(u_t)$, or should both nodes be migrated to a new cluster? Moreover, an algorithm may be required to pro-actively migrate (resp. swap) additional nodes.

C) *Which nodes to evict?* There may not exist sufficient space in the desired destination cluster. In this case, the algorithm needs to decide which nodes to evict, to free up space.

Our Contributions. This paper introduces the online balanced repartitioning problem. We consider deterministic algorithms and make the following technical contributions:

1. *Online Rematching:* For the special case of online rematching ($k = 2$, but arbitrary ℓ), Theorem 1 presents a greedy online algorithm which is almost optimal: it is 7-competitive and we prove a lower bound of 3.
2. *Lower Bounds:* While in a setting without augmentation, a $k-1$ lower bound for the competitive ratio of any online algorithm follows from a simulation of online paging, in Theorem 2, we show a lower bound which is strictly larger than k, for any $\alpha > 0$. Intriguingly, we show that the online repartitioning problem remains hard even with augmentation. In particular, in Theorem 3 we prove that no augmented online algorithm can achieve a competitive ratio below k, as long as it cannot solve the problem trivially by placing all nodes into a single cluster. In contrast, online paging is known to become constant competitive with constant augmentation [28].
3. *Online Balanced Repartitioning:* Our main technical contribution stated in Theorem 4 is a non-trivial $O(k \log k)$-competitive algorithm for the setting with 4-augmentation.

Observe that none of our bounds depends on ℓ. This is interesting, as for example, in our motivating virtual machine collocation problem, k is likely to be small: a server typically hosts a small number of virtual machines (e.g., related to the constant number of cores on the server).

Paper Organization. The remainder of this paper is organized as follows. After reviewing related work in Section 2, we start by discussing the special case of matchings ($k = 2$) in Section 3. We consider lower bounds for the general setting with and without augmentation in Section 4. Section 5 is then devoted to the presentation and analysis of CREP, an augmented deterministic algorithm. We conclude in Section 6. To simplify presentation, we postpone some technical details to the Appendix.

2 Related Work

Our work assumes a new perspective on several classic algorithmic problems. The static version of our problem (without migration) is the minimum balanced graph partitioning problem (where ℓ is the number of components). This problem is known to be NP-complete, and cannot even be approximated within any finite factor [4]. The static variant where $k = 2$ corresponds to a maximum matching problem, which is polynomial-time solvable. The static variant where $\ell = 2$ corresponds to the minimum bisection problem [16] and can be approximated within a factor of $O(\log^{1.5} n)$ from the minimum cost [19]. The concept of cluster-size augmentation is inspired by offline *bi-criteria* approximations to graph partitioning, in particular the (ℓ, δ)-balanced graph partitioning problem [4], where the graph has to be partitioned in ℓ components of size less than $\delta \cdot \frac{n}{\ell}$, as well as by the concept of c-balanced cuts used by Arora, Rao, and Vazirani [5], where both partitioned components should be of size at least $c \cdot n$.

In terms of online algorithms, the subproblem of finding a good trade-off between serving requests remotely (at a low but repeated communication cost) or migrating nodes together (entailing a potentially high one-time cost α), is essentially a ski rental or rent-or-buy problem [20,21]. A similar tradeoff also arises in the context of online page and server migration problems [8,9], where requests appear in a metric space [10] or graph [6] over time, and need to be served by one [9] or multiple [18] servers. However, in BRP, the number of possible node-cluster configurations is large, rendering it difficult to cast the problem into an online metrical task system. Moreover, in contrast to most online migration problems, which typically optimize the placement of a page or server with respect to the request locations, in our model, *both* end-points of a communication request are subject to optimization. A second difference to the usual models studied in the literature (where requests appear at specific locations in the metric space) is that in our problem a request only reveals partial and binary information about the optimal location (resp. configuration) to serve it: the request can be served at cost zero whenever the communicating are collocated.

Our model can be seen as a generalization of online paging [17,22,23,29], sometimes also referred to as online caching, where requests for data items arrive over time and need to be served from a cache of finite capacity, and where the number of cache misses must be minimized. The online caching and paging problem was first analyzed in the framework of the competitive analysis by Sleator and Tarjan [28], who presented a $kALG/(kALG - kOPT + 1)$-competitive algorithm,

where $kALG$ is the cache size of the online and $kOPT$ the cache size of the offline algorithm. The authors also proved that no deterministic online algorithm can beat this ratio. In the classic caching model and its variants [11,12,23], items need to be put into the cache upon each request, and the problem usually boils down to finding a smart eviction strategy, such as Least Recently Used (LRU) or Flush-When-Full (FWF). In contrast, in our setting, requests can be served remotely. In this light, our model is reminiscent of caching models *with bypassing* [1,15]. In fact, it is easy to see that in a scenario with $\ell = 2$ clusters, online paging can be simulated (cf. Section A.5): in this simulation, one cluster can be used as the cache and the other cluster as the slow memory; by the corresponding problem reduction we also obtain a $k - 1$ lower bound for our problem without augmentation. However, in general, in our model, the "cache" is *distributed*: requests occur *between* nodes and *not to* nodes, and costs can be saved by collocation.

BRP also has connections to online packing problems, where items of different sizes arriving over time need to be packed into a minimal number of bins [25,27]. In contrast to these problems, however, in our case the objective is not to minimize the number of bins but rather the number of "links" between bins, given a fixed number of bins.

Finally, our model also connects to recent work on online clique and correlation clustering [3,7,14,26]. In this prior work, nodes and/or links can appear over time, but the underlying communication graph remains invariant.

3 Online Rematching

Let us first consider the special case where clusters are of size two ($k = 2$, arbitrary ℓ). This is essentially an online maximal (re)matching problem: clusters of size two host (resp. "match") exactly one pair of nodes, and maximizing pair-wise communication within each cluster is equivalent to minimizing inter-cluster communication. In a $k = 2$ scenario, the question of which node to evict is trivial: there is simply no choice. The problem can also be seen from a ski-rental perspective: one has to identify a good tradeoff between serving requests remotely ("renting") and migrating the communicating nodes together ("buy").

A natural greedy online algorithm Greedy to solve this problem proceeds as follows: For each cluster C_i, hosting nodes u_i, v_i, we count the total number of inter-cluster requests over time for its nodes. After 3α requests occur between nodes inside any cluster C_1 to nodes outside the cluster, we identify the cluster C_2 with which C_1 communicated most frequently in this time period. We then collocate u_1 with the single node in C_2 (v_2 or u_2) with which it communicated the most—ties broken arbitrarily and without involving any other clusters in the repartitioning. Afterwards, we reset all pairwise communication counters involving nodes from (old) clusters C_1 and C_2 (i.e., u_1, u_2, v_1, v_2).

We can prove the following (see Appendix):

Theorem 1. Greedy *is* 7*-competitive. No deterministic online algorithm achieves a competitive ratio below* 3 *when* $|\sigma| \to \infty$.

4 Lower Bounds for Online Balanced Repartitioning

Our problem is generally hard to approximate online. While it is easy to see that a lower bound of $k - 1$ follows by simulating online paging (ufsing only two servers), in the following we prove a strictly larger lower bound (cf. Theorem 2). In fact, we observe that, even with augmentation, our problem is hard to approximate online: as long as the augmentation is less than what would be required to solve the partitioning problem trivially, by putting all nodes into the same cluster (i. e., $\delta < \ell$), no deterministic online algorithm can achieve a competitive ratio better than k (cf. Theorem 3). This highlights an intriguing difference from online paging, where the competitive ratio becomes constant under augmentation [28]. Our lower bounds are independent of the initial configuration: both Off and On start off having the nodes placed identically in clusters.

Theorem 2. *No deterministic online algorithm can achieve a competitive ratio smaller than* $k + \frac{k-2}{2\alpha}$*, independently of* ℓ.

Interestingly, an adversary can outwit any online algorithm, even in the setting with augmentation. In the following, we consider online algorithms which, compared to Off, have δ-times more space in each cluster.

Theorem 3. *No* δ*-augmented deterministic online algorithm can achieve a competitive ratio smaller than* k*, as long as* $\delta < \ell$.

5 CREP Algorithm: An $O(k \log k)$ Upper Bound

The main technical contribution of this paper is an online *Component-based REPartitioning algorithm (*CREP*)* which achieves an almost tight upper bound matching the k lower bound of Theorem 3 with augmentation at least 4. Intuitively, it helps to think of a 4-augmented algorithm as one that can use twice as many clusters, each having twice as much space (though this is a special case of the definition of augmentation). Formally, we claim:

Theorem 4. *With augmentation at least 4,* CREP *is $O(k \log k)$-competitive.*

CREP is summarized in Algorithm 1. The algorithm is non-trivial and relies on the following basic ideas:

1. *Communication components.* CREP groups nodes which have recently communicated into *components*. Once the cumulative communication cost of a group of nodes distributed across two or more components exceeds a certain threshold, CREP merges them into a single component, by collocating them in the same cluster. That is, we maintain a logical, time-varying weighted component graph $G_t = (\Phi_t, E_t, w_t)$, where Φ_t is the set of components immediately after request t has been issued, the edges E_t connect components which communicated at least once during this epoch, and w_t is the number of communication requests between the corresponding two nodes in this epoch. In other words, an edge $(i, j) \in E$ between two components $\phi_i, \phi_j \in \Phi_t$ indicates that the two components (resp. the corresponding nodes in ϕ_i and those in ϕ_j) were involved in $w_{ij} > 0$ requests. Although the graph G_t changes over time (when components are merged or split according to CREP), when the time is clear from the context, we drop the time-index and simply write $G = (\Phi, E, w)$. Edges disappear (and their weights are reset) when the components are merged. For a *component set* $X = \{\phi_i, \phi_j, \dots\} \subseteq \Phi$, let $|X|$ denote the number of components in X. We call $\mathrm{vol}(X) = \sum_{\phi \in X} |c|$ the *volume* of the set and $\mathrm{com}(X) = \sum_{\phi_i, \phi_j \in X} w_{ij}$ the *communication cost* among the members of X.
2. *Component epochs.* We analyze CREP in terms of component-wise epochs. A component epoch starts with the first request between two individual nodes (singleton components), and ends when: (i) the size of the component (the number of nodes in the component) exceeds k and (ii) the accumulated communication between the components exceeds a certain threshold. CREP maintains the invariant that components are never split during an epoch, that is, once two nodes of the same component epoch are placed together in a cluster, they will remain in the same cluster in the remainder of the epoch (but they may possibly be migrated together to a new cluster). As such, when a component set X is merged into a new component (Line 7), CREP tries to migrate all the components to the cluster of the largest component (ties broken arbitrarily). If there is not enough reserved space in the cluster, then all components are migrated to a new cluster. If on the other hand $\mathrm{vol}(X)$ exceeds k, the component-epoch ends, and all $\phi \in X$ are reset to singleton components (Line 19). More specifically, according to Algorithm 1, *two* termination criteria have to be fulfilled for a component set Y to end an epoch: $\mathrm{vol}(Y) > k$ *and* $\mathrm{com}(Y) \geq \mathrm{vol}(Y) \cdot \alpha$. This non-trivial criterion is critical, to keep the migration cost competitive. An epoch that ends as a result of a set Y is referred to as a Y-epoch.
3. *Space reservations.* In order to keep the number of migrations low, CREP performs space reservations in clusters. Whenever CREP migrates a component ϕ into a cluster, it reserves additional space $\mathrm{reserve}(\phi) = \min\{k - |\phi|, |\phi|\}$. As we will prove, these proactive space reservations can ensure that a component has to be migrated again only after its size doubles. For a cluster s, let $\mathrm{reserved}(s), \mathrm{occupied}(s)$ and $\mathrm{spare}(s)$ denote the reserved, occupied and spare (unreserved) space in s, where always $\mathrm{reserved}(s) + \mathrm{occupied}(s) + \mathrm{spare}(s) = 2k$. Similarly, for a component ϕ let $\mathrm{reserved}(\phi)$ denote the amount of its reserved space that is still available in its current cluster.

The remainder of this section is devoted to the proof of Theorem 4. The proof unfolds in a number of observations and lemmas. We first observe, in Property 1, that indeed, it is always possible to find a cluster where the to-be-merged components fit. We then derive an upper bound on CREP's cost per component epoch and a lower bound on the optimal offline cost per component epoch. Finally, in the Appendix, we show that the competitive ratio is also bounded with respect to incomplete epochs.

We start by observing that there always exists a cluster which can host the entire merged component, including the required reserved space without any evacuation, i.e., its spare space is at least k.

Algorithm 1 CREP with 4 Augmentation

1: Construct graph $G = (\Psi, E, w)$ with singleton components: one component per node. Set $w_{ij} = 0$ for all $\{v_i, v_j\} \in \binom{V}{2}$. For each component ϕ_i, reserve space reserve$(\phi_i) = 1$.

2: **for** each new request $\{u_t, v_t\}$ **do**

▷ Keep track of communication cost.

3: Let $\phi_i = \Phi(u_t)$ and $\phi_j = \Phi(v_t)$ be the two components that communicated.

4: **if** $\phi_i \neq \phi_j$ **then**

5: $w_{ij} \leftarrow w_{ij} + 1$

6: **end if**

▷ Merge components.

7: Let X be the largest cardinality set with $\text{vol}(X) \leq k$ and $\text{com}(X) \geq (|X| - 1) \cdot \alpha$

8: **if** $|X| > 1$ **then**

9: Let $\phi_0 = \bigcup_{\phi_i \in X} \phi_i$ and for all $\phi_j \in \Phi \backslash X$ set $w_{0j} = \sum_{\phi_i \in X} w_{ij}$.

10: Let $\phi \in X$ be the component having the largest reserved space.

11: **if** reserved$(\phi) \geq \text{vol}(X) - |\phi|$ **then**

12: Migrate ϕ_0 to the cluster hosting ϕ

13: Update reserved$(\phi_0) =$ reserved$(\phi) - (\text{vol}(X) - |\phi|)$

14: **else**

15: Migrate ϕ_0 to a cluster s with spare$(s) \geq \min(k, 2|\phi_0|)$

16: Set reserved$(\phi_0) = \min(k - |\phi_0|, |\phi_0|)$

17: **end if**

18: **end if**

▷ End of a Y-epoch.

19: Let Y be the smallest components set with $\text{vol}(Y) > k$ and $\text{com}(Y) \geq \text{vol}(Y) \cdot \alpha$

20: **if** $Y \neq \emptyset$ **then**

21: Split every $\phi_i \in Y$ into ϕ_i singleton components and reset the weights of all edges involving at least one newly created component. Reserve one additional space for each newly created component. If necessary, migrate at most $\text{vol}(Y)/2 + 1$ singletons to clusters with spare space.

22: **end if**

23: **end for**

Property 1. At any point in time, a cluster exists having at least k spare space.

So indeed, CREP can always place a merged component greedily into clusters—no global component rearrangement is necessary. On the other hand, augmenting the cluster size allows CREP to reserve additional space for migrated components. As we show in the following, this ensures that each node is migrated at most $\log k$ times (rather than k) during the formation of a component.

Upper bound on CREP's costs. The online algorithm's cost during each epoch consists of the *communication cost*, which amounts to the number of communication requests that were served remotely, and the *migration cost*, which is equal to the number of node migrations. The following properties provide upper bounds for both kinds of costs for a single component:

Property 2. At any point in time, consider a component c induced by the communication pattern in this epoch, then:

1. The communication cost between nodes in ϕ is, in this epoch, at most $(|\phi| - 1) \cdot \alpha$.
2. The migration cost of nodes in ϕ is, in this epoch, at most $(|\phi| \log |\phi|) \cdot \alpha$.

Proof (Proof of Property 2). The two properties are proved in turn.

Property 2.1. We prove this property by induction on the *merging sequence*, i. e., the sequence of merges that includes all the nodes in ϕ from the time when they were singletons, ordered by time. To establish the base case, consider the first merge of nodes in ϕ, where X was a set of singletons (Line 7) and $|X|$ singleton components were combined into a new component $\phi_0 = \cup_{\phi_i \in X} \phi_i$. By CREP's merging condition, the cost up to this point is equal to $(|X| - 1) \cdot \alpha = (|\phi_0| - 1) \cdot \alpha$. For the inductive step, consider again that X' is merged into ϕ_0 and suppose that the communication cost paid for each component $\phi' \in X'$ is $(|\phi'| - 1) \cdot \alpha$. After the merge, CREP's total communication cost is equal to $(|X'| - 1) \cdot \alpha + \sum_{\phi_i \in X'} (|\phi_i| - 1) \cdot \alpha = \left(\sum_{\phi_i \in X'} |\phi_i| \right) \cdot \alpha - \alpha = (|\phi_0| - 1) \cdot \alpha$ and the induction holds.

Property 2.2. First observe that any node u which belongs to ϕ is migrated at most $\log |\phi|$ times during an epoch. To see this, suppose that u was just migrated into a cluster and that the size of u's

current component is $|\phi'|$. From Property 1, we know that u will not be migrated as a consequence of a merge that does not involve u's component (i.e., it will never be evicted). Furthermore, due to the existence of reserved space, u will stay in the same cluster as long as the size of its current component ϕ' remains smaller or equal to $2|\phi'|$. Since the size of u's component between any consecutive migrations doubles, the total migrations can be at most $\log |\phi|$. This implies the total number of migrations pertaining to all nodes in ϕ is at most $|\phi| \log |\phi|$.

Using Property 2, we can bound the migration and communication cost of a Y-epoch:

Lemma 1. *Consider the end of a Y-epoch (Line 19).* CREP *migrates at most $\sum_{\phi_i \in Y} |\phi_i| \log |\phi_i| \leq vol(Y) \cdot \log k$ nodes and serves at most $2\, vol(Y) \cdot \alpha$ remote requests during this epoch for nodes in Y.*

The proof of Lemma 1 follows directly from Property 2 and is omitted.

Lower bound on Off's cost. Having derived an upper bound on CREP's cost, we next compute a lower bound of Off's cost.

Lemma 2. *By the end of a Y-epoch,* Off *pays at least $vol(Y)/k \cdot \alpha$ communication cost (during this epoch) for nodes in Y.*

To establish the above lower bound, we will need two useful properties.

Property 3. Consider any component ϕ in the current epoch and *any* partition of ϕ into two non-empty disjoint sets B and W, with $B \cup W = \phi$. During the creation of ϕ (by merging), there were at least α communication requests between nodes in B and W.

Proof (Proof of Property 3). Consider the tree T which describes how component ϕ merged from singletons into ϕ during the current epoch. The leafs of the tree are the nodes in ϕ and each internal node corresponds to a component set X that was found in Line 7 of the algorithm, and entails a merge to a new component ϕ_0. Now color the leafs of the tree according to W and B. Since both sets are non-empty there must exist an internal node τ in T, whose descendant leafs in the subtree are not colored identically. Let $B' = \{b_1, b_2, \ldots, b_p\}$ be the child components of τ which are in B, and let $W' = \{w_1, w_2, \ldots, w_q\}$ be the components which are children of τ and which are in W. Let X be the set corresponding to the descendant leafs of τ and note that $X = B' \cup W'$ and $|X| = p+q$. Since neither B' nor W' were merged earlier, the total communication cost among B' (resp. W') is at most $(p-1)\alpha$ (resp. $(q-1)\alpha$), which sums up to at most $(p+q-2)\alpha$. But since X is witnessing a communication cost of at least $(|X|-1)\alpha = (p+q-1)\alpha$, there must have occurred at least $(p+q-1)\alpha - (p+q-2)\alpha = \alpha$ communication cost between the nodes in B and the nodes in W during the current epoch.

Property 4. Consider *any* two non-empty disjoint subsets of components $U, V \subset Y$ s.t. $U \cup V = Y$ and $\text{vol}(U) < k$ (i.e., U's components fit in one cluster). The inter-communication cost between U and V is at least α during the Y-epoch.

Proof (Proof of Property 4). The proof follows from the minimality of Y and because no merge involving components happened in Y. From the Y-epoch termination condition, we have $\text{com}(Y) = \text{com}(U) + \text{com}(V) + \text{inter}(U, V) \geq \text{vol}(Y) \cdot \alpha$, where $\text{inter}(U, V)$ is the inter-communication cost between U and V. Assume for the sake of contradiction that $\text{inter}(U, V) < \alpha$. Recall that $\text{vol}(U) \leq k$ and, since the components in U have not been merged yet, $\text{com}(U) \leq (|U|-1) \cdot \alpha \leq (\text{vol}(U)-1) \cdot \alpha$ (the equality holds when $|U| = 1$). We therefore have

$$\begin{aligned} \text{com}(V) &= \text{com}(Y) - \text{com}(U) - \text{inter}(U, V) \\ &> \text{vol}(Y) \cdot \alpha - (\text{vol}(U) - 1) \cdot \alpha - \alpha = \text{vol}(V) \cdot \alpha. \end{aligned}$$

We obtain the desired contradiction by distinguishing between two cases: (i) If $\text{vol}(V) > k$, then $V \subset Y$ meets both termination conditions of a component epoch (Line 19), and thus the minimality of set Y is contradicted. (ii) Next, consider that $\text{vol}(V) \leq k$ and notice that it must hold that $|V| > 1$ (otherwise $\text{com}(V) = 0$). Since the components in V have not been merged yet, $\text{com}(V) \leq (|V|-1) \cdot \alpha \leq (\text{vol}(V)-1) \cdot \alpha \leq \text{vol}(V) \cdot \alpha$, which is again a contradiction.

Proof (Proof of Lemma 2). We now use Properties 3 and 4 to lower bound Off's cost. First, it follows from Property 3 that Off cannot gain by splitting any component $\phi \in Y$ between different clusters (which would only increase its cost). The question is then, how much can Off reduce the

inter-communication cost by arranging $\phi \in Y$ more efficiently? Let R_{intra} be the number of inter-communication requests (of CREP) OFF did not pay by placing the components in an optimal way. Furthermore, denote by s the number of clusters OFF used and by $b_j \leq k$ the number of nodes OFF placed in each cluster $j = 1, \ldots, s$. Consider any cluster j. In the best case for OFF, each of the b_j nodes in the cluster is a singleton component which CREP placed in a different cluster. It follows that OFF's saved cost cannot be greater than $(b_j - 1)\alpha$: otherwise CREP would have merged the b_j components into a single component. Observe also that, although OFF aims to put the components in as few possible clusters, by a simple pigeonhole argument, s must be at least $\mathrm{vol}(Y)/k$. Combining the above, the number of requests R_{inter} OFF serves remotely (assuming that no node was migrated during the Y-epoch) is

$$R_{inter} = \mathrm{vol}(Y) \cdot \alpha - R_{intra} \geq \mathrm{vol}(Y) \cdot \alpha - \sum_{j=1}^{s} (b_j - 1) \cdot \alpha$$
$$= \left(\mathrm{vol}(Y) - \sum_{j=1}^{s} b_j + \sum_{j=1}^{s} 1 \right) \cdot \alpha \geq \sum_{j=1}^{\mathrm{vol}(Y)/k} 1 = \frac{\mathrm{vol}(Y)}{k} \cdot \alpha,$$

where the last step follows from the fact that $\sum_{j=1}^{s} b_j = \mathrm{vol}(Y)$.

It remains to show that OFF cannot decrease R_{inter} any further by swapping nodes during the Y-epoch. From Property 4, the nodes in each cluster j communicated at least α times with clusters $i \neq j$. Since any swap that OFF performs between two clusters costs at least 2α, a swap between the involved clusters can only be beneficial to ON. Considering that there are at least $s \geq \mathrm{vol}(Y)/k$ clusters, even with migrations, OFF's cost will be at least $\mathrm{vol}(Y)/k \cdot \alpha$.

Incomplete Component Epoch. So far, we have quantified the cost that CREP and OFF pay *at the end of each epoch*. It remains to account for the costs that CREP accumulates in incomplete epochs.

First, let us observe that the edge weights w of incomplete epochs in the component graph are naturally bounded: at some point, the edge will cause a merge, or end the epoch. By dividing the edges of the component graph G into *light edges* and *heavy edges*, we can claim the following:

Property 5. For every edge (ϕ_i, ϕ_j) in the component graph, at any given time:

1. If $|\phi_i| + |\phi_j| \leq k$ (we call this a *light edge*), the edge has cost at most α.
2. If $|\phi_i| + |\phi_j| > k$ (we call this a *heavy edge*), the edge cost is at most $(|\phi_i| + |\phi_j|)\alpha \leq (2k)\alpha$

The claim is implied by the definition of CREP. In the first case, if the (ϕ_i, ϕ_j) edge cost was larger than α, CREP would have merged ϕ_i and ϕ_j into a new component. Similarly, in the second case, if the edge (ϕ_i, ϕ_j) cost was larger than $(|\phi_i| + |\phi_j|)\alpha$, CREP would have ended the epoch.

Let us consider the request sequence σ at some time t. Recall that at the end of an epoch, we reset all involved edge weights, and charge OFF for them. So at time t, we have not taken into account yet the communication requests that were not reset. For any two nodes u and v, we consider all their communication requests since the last time they belonged to the same Y, at the end of a Y-epoch. All these requests belong to what we call the *last epoch*. Note that σ may not contain any complete epochs at all. But every request $\{u, v\} \in \sigma$ must belong to some Y-epoch or to the last epoch. Using Property 5 we can claim the following:

Lemma 3. *The competitive ratio of* CREP *for communication requests which belong to the* last epoch*, is bounded by* $O(k \log k)$.

The competitive ratio of CREP follows from Lemmas 1, 2, and 3.

6 Conclusion

This paper initiated the study of a natural dynamic partitioning problem which finds applications, e.g., in the context of virtualized distributed systems subject to changing communication patterns.

We derived different upper and lower bounds, both for the general case as well as for a special case describing a matching problem. While the derived competitive ratios are sometimes linear or even super-linear in k, they do not depend on ℓ: We believe that this is attractive in practice: for

example, while the number of servers in a datacenter (i.e., ℓ) can be large, the number of virtual machines hosted per server (e.g., the number of cores) is usually small.

The main open question raised by our work regards the optimality of our upper bound: currently, the upper and lower bounds are off by a logarithmic factor. Moreover, it will be interesting to explore randomized settings: While we have some early positive results on the potential of randomization for special problem instances, the feasibility of $o(k)$-competitive randomized algorithms remains an open problem.

Acknowledgements. We thank Marcin Bienkowski for many inputs and feedback on this paper.

References

1. A. Adamaszek, A. Czumaj, M. Englert, and H. Räcke. An $O(\log k)$-competitive algorithm for generalized caching. In *Proc. 23rd SODA*, pages 1681–1689, 2012.
2. M. Al-Fares, A. Loukissas, and A. Vahdat. A scalable, commodity data center network architecture. *ACM SIGCOMM Computer Communication Review*, 38(4):63–74, 2008.
3. K. Andreev and H. Räcke. Balanced graph partitioning. In *Proc. 16th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2004.
4. K. Andreev and H. Räcke. Balanced graph partitioning. *Theory of Computing Systems*, 39(6):929–939, 2006.
5. S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)*, 56(2):5, 2009.
6. B. Awerbuch, Y. Bartal, and A. Fiat. Competitive distributed file allocation. *Inf. Comput.*, 185(1):1–40, Aug. 2003.
7. N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Mach. Learn.*, 56(1-3):89–113, June 2004.
8. Y. Bartal, M. Charikar, and P. Indyk. On page migration and other relaxed task systems. *Theoretical Computer Science*, 268(1):43–66, 2001. Also appeared in *Proc. of the 8th SODA*, pages 43–52, 1997.
9. M. Bienkowski, A. Feldmann, J. Grassler, G. Schaffrath, and S. Schmid. The wide-area virtual service migration problem: A competitive analysis approach. *IEEE/ACM Transactions on Networking*, 22(1):165–178, 2014.
10. A. Borodin, N. Linial, and M. E. Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM*, 39(4):745–763, 1992. Also appeared in *Proc. of the 19th STOC*, pages 373–382, 1987.
11. M. Brehob, R. J. Enbody, E. Torng, and S. Wagner. On-line restricted caching. *Journal of Scheduling*, 6(2):149–166, 2003.
12. N. Buchbinder, S. Chen, and J. Naor. Competitive algorithms for restricted caching and matroid caching. In *Proc. of the 22thEuropean Symp. on Algorithms (ESA)*, pages 209–221, 2014.
13. M. Chowdhury, M. Zaharia, J. Ma, M. I. Jordan, and I. Stoica. Managing data transfers in computer clusters with orchestra. *SIGCOMM Comput. Commun. Rev.*, 41(4):98–109, 2011.
14. C. H. Q. Ding, X. He, H. Zha, M. Gu, and H. D. Simon. A min-max cut algorithm for graph partitioning and data clustering. In *Proc. IEEE International Conference on Data Mining (ICDM)*, pages 107–114, 2001.
15. L. Epstein, C. Imreh, A. Levin, and J. Nagy-György. Online file caching with rejection penalties. *Algorithmica*, 71(2):279–306, 2015.
16. U. Feige and R. Krauthgamer. A polylogarithmic approximation of the minimum bisection. *SIAM Journal on Computing*, 31(4):1090–1118, 2002.
17. A. Fiat, R. M. Karp, M. Luby, L. A. McGeoch, D. D. Sleator, and N. E. Young. Competitive paging algorithms. *Journal of Algorithms*, 12(4):685–699, 1991.
18. A. Fiat, Y. Rabani, and Y. Ravid. Competitive k-server algorithms. *J. Comput. Syst. Sci.*, 48(3):410–428, 1994.
19. R. Krauthgamer and U. Feige. A polylogarithmic approximation of the minimum bisection. *SIAM Review*, 48(1):99–130, 2006.
20. A. Kumar, A. Gupta, and T. Roughgarden. A constant-factor approximation algorithm for the multicommodity rent-or-buy problem. In *Proc. 43rd Symposium on Foundations of Computer Science (FOCS)*, 2002.
21. Z. Lotker, B. Patt-Shamir, and D. Rawitz. Rent, lease or buy: Randomized algorithms for multislope ski rental. In *Proc. of the 25thSymp. on Theoretical Aspects of Computer Science (STACS)*, pages 503–514, 2008.
22. L. A. McGeoch and D. D. Sleator. A strongly competitive randomized paging algorithm. *Algorithmica*, 6(6):816–825, 1991.
23. M. Mendel and S. S. Seiden. Online companion caching. *Theoretical Computer Science*, 324(2–3):183–200, 2004.
24. J. C. Mogul and L. Popa. What we talk about when we talk about cloud network performance. *ACM SIGCOMM Computer Communication Review*, 42(5):44–48, 2012.
25. P. V. Ramanan, D. J. Brown, C. C. Lee, and D. T. Lee. On-line bin packing in linear time. *Journal of Algorithms*, 10(3):305–326, 1989.

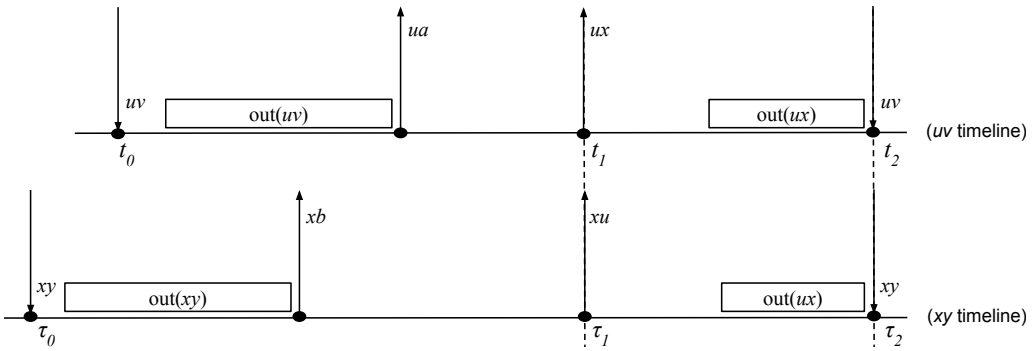

Fig. 1. The illustration, which accompanies the proof of Lemma 4, shows a sequence of events leading GREEDY to merge nodes u and v, as well as x and y, at time $t_2 = \tau_2$. The horizontal axis denotes time. Each of the two timelines corresponds to the swaps and requests involving a pair (uv or xy). Arrows correspond to merges (pointing down) and splits (pointing up) of node pairs by GREEDY: for example, the down-pointing arrow on the top right indicates that GREEDY collocated at time t_2 nodes u and v. Boxes correspond to requests GREEDY served remotely. For instance, the events on the right-most side correspond to the sequence of $\lambda\alpha$ requests that cause GREEDY to swap at t_2.

26. S. E. Schaeffer. Graph clustering. *Computer Science Review*, 1(1):27 – 64, 2007.
27. S. S. Seiden. On the online bin packing problem. *Journal of the ACM*, 49(5):640–671, 2002.
28. D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.
29. N. E. Young. On-line caching as cache size varies. In *Proc. of the2ndACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 241–250, 1991.

A Deferred Proofs

A.1 Proof of Theorem 1

We prove the theorem by showing the following two lemmas.

Lemma 4. GREEDY *is* 7*-competitive.*

Proof. We analyze the competitive ratio of GREEDY by considering a parametrized algorithm GREEDY(λ), which acts exactly like GREEDY, but using a threshold $\lambda\alpha$. To prove our claim, we consider each swap that GREEDY(λ) performs, and show that OFF must have paid at least a λ-dependent fraction of the online algorithm's cost (counting all communication and migration costs in the period preceding the swap and considering all nodes involved). After also taking remaining costs into account, we derive the competitive ratio by minimizing over λ. More concretely, suppose that at time $t_2 = \tau_2$, GREEDY(λ) merged nodes uv and xy, where, at time t_1, the partitioning was ux and vy—see Figure 1.

Greedy(λ). We proceed to account for all costs of the online algorithm pertaining to the swap, i. e., the communication and migration costs pertaining to pair uv starting from the last time the pair was merged by GREEDY(λ), denoted as t_0, and until t_2, as well as the respective costs of pair xy in the corresponding period $(\tau_0, \tau_2]$. If t_2 (or τ_2) is the first time uv (or xy) were collocated, then set $t_0 = 0$ (or respectively $\tau_2 = 0$). Obviously, since GREEDY(λ) only merged the pairs once in the considered period, it paid exactly 2α migration cost. Furthermore, if t_1 is the last time u was merged with some other node prior to t_2, this node must be x, meaning that during period $(\tau_1 = t_1, \tau_2 = t_2]$, GREEDY($\lambda$) served $\lambda\alpha$ requests remotely between the cluster C_1 holding u and x, and other clusters (boxes out(ux) in the figure). W.l.o.g., let C_1 be the cluster that initiated the swap (its outgoing request-counter reached $\lambda\alpha$). Algorithm GREEDY(λ) might have also served $\beta\lambda\alpha$ requests between C_2 and some other clusters (not including C_1), where $0 \leq \beta \leq 1$ will be bounded later. Summarizing, the online algorithm's cost is GREEDY(λ) $< (2 + (1 + \beta)\lambda)\alpha$.

Off. W.l.o.g., let us focus on u and v. Depending of whether OFF served the out(ux) requests remotely or not, we distinguish two cases:

1. **Off had u and v collocated during parts of $(t_0, t_2]$.** As we show next, independently of x and y, OFF must have paid at least $\min\{\alpha, \lambda\alpha/2\}$. To see this, notice that OFF had two options: (a) *First,* OFF *merged u and v at some instant during* (t_0, t_2). Splitting the 2α

migration cost Off paid for the four involved nodes, and charging the two quarters pertaining to nodes u and v, we find that in this case Off $\geq \alpha$. (b) *Second,* Off *merged u and v before t_0.* Obviously, this can only happen if $t_0 > 0$. In this case, Off must have paid $\lambda\alpha$ for the outgoing requests out(uv) from the cluster hosting u and v between $(t_0, t_2]$. By distributing again the costs over the four involved nodes and charging the two quarters pertaining to nodes u and v, it follows that Off $\geq \lambda\alpha/2$.

2. **Off kept u and v in different clusters during $(t_0, t_2]$.** In contrast to the previous case, here Off has the option to distribute its communication costs arbitrarily amongst u and v, or x and y. We therefore need to examine Off's actions w. r. t. to x and y, as well as u and v. We distinguish three sub-cases: (a) *First,* Off *merged x and y at some instant during (τ_0, τ_2).* This is the symmetric case to Case 1(a). As before, we can charge Off $= \alpha$ migration cost (we distribute the 2α cost paid for the four involved nodes). (b) *Second,* Off *merged x and y before τ_0.* This is the symmetric case to Case 1(b). As before, after distributing the cost among all involved nodes, we find that Off $\geq \lambda\alpha/2$ because of the outgoing requests from x and y during $(\tau_0, \tau_2]$ (box out(xy) in the figure). So for cases (a) and (b), Off $\geq \min\{\alpha, \lambda\alpha/2\}$. (c) *Third,* Off *kept x and y at different clusters during $(\tau_0, \tau_2]$.* By the swapping condition of Greedy(λ), we know that during $(t_1, t_2]$, a total of out$(ux) = \lambda\alpha$ requests were served between cluster C_1 hosting nodes u, x and all other clusters. In addition, there may have been $\beta\lambda\alpha$ requests between cluster C_2 and some other clusters (for instance cluster C''). There are three types of requests: R_{12} requests between clusters C_1 and C_2, R_1 requests involving cluster C_1 but not C_2, and $R_2 = \beta\lambda\alpha$ requests involving cluster C_2 but not C_1. Furthermore, $R_{12} + R_1 = \lambda\alpha$. Now, since Off kept both u and v, as well as x and y in different clusters, it must have paid for all R_{12}. Regarding R_1, since Off performed no swaps during $(\tau_0, \tau_2]$, it was only able to save $R' \leq R_1$ requests by smartly placing u and x with two other nodes p and q at some earlier time. Yet, since the online algorithm always collocates the most heavily communicating nodes and additionally the saved requests R' can only correspond to two node pairs (up and xq), it must hold also that $2R' \leq R_{12}$, so $R' \leq 2\lambda\alpha/3$ and Off pays at least $\lambda\alpha - R' \geq \lambda\alpha/3$. Having appropriately placed the nodes between clusters C_2 and C'', Off can avoid paying for the $\beta\lambda\alpha$ requests.

We proceed to argue about $\beta = R_2/(\lambda\alpha)$. Let us begin with some useful inequalities: $R_1 + R_{12} = \lambda\alpha$ (since cluster C_1 triggered the swap), $R_2 + R_{12} < \lambda\alpha$ (otherwise C_2 would have triggered the swap), and $R' < 2R_{12}$ (since Greedy(λ) collocates the heaviest communicating pairs). Off has two options: either to save the R_{12} requests paying at least $\lambda\alpha - R_{12}$, or to save R' out of the R_1 requests, paying at least $\lambda\alpha - R'$. Since $R' \leq 2R_{12}$, in both cases Off's cost is larger than $\lambda\alpha - 2R_{12}$. Using the inequality $R_{12} < \lambda\alpha - R_2 = (1-\beta)\lambda\alpha$, we therefore find that, no matter how the offline algorithm partitions the nodes, it must pay Off $\geq \lambda\alpha - 2R_{12} > (2\beta - 1)\lambda\alpha$.

Remaining costs. The only costs not accounted for by the previous analysis correspond to requests between node pairs at the end of the sequence that (i) did not not cause Greedy(λ) to swap (they were less than $\lambda\alpha$), and (ii) were not reset during some previous swap. Suppose for instance that, just before the end of the request sequence, the adversary issues R requests between u and v. We claim that for $R < \lambda\alpha$ the competitive ratio cannot be larger than when $R = \lambda\alpha$ (i.e., the case considered by our previous analysis): no matter how Off and Greedy(λ) have arranged u and v when the R requests are issued, Off will always pay any induced costs either prior (if Off migrates) or at the same time (if Off served the requests remotely) as the online algorithm. Therefore, issuing additional $\lambda\alpha - R$ requests and forcing Greedy(λ) algorithm to swap can only increase the offline algorithm's cost as much as the online cost is increased, and the claim follows.

Optimizing over λ, we obtain Greedy's competitive ratio

$$\rho(\text{Greedy}) = \min_{\lambda>0} \max\left\{ \frac{(2+(1+\beta)\lambda)\,\alpha}{\min\{\lambda\alpha/3, \alpha\}}, \frac{(2+(1+\beta)\lambda)\,\alpha}{(2\beta-1)\lambda\alpha} \right\} < \min_{\lambda>0} \frac{2+(1+2/3)\lambda}{\min\{\lambda/3, 1\}} = 7,$$

which is achieved for $\lambda^* = 3$.

This is not far from the optimal, since no deterministic online algorithm achieves a competitive ratio below 3.

Lemma 5. *For $k = 2$ and any request sequence σ of length $|\sigma|$, no deterministic online algorithm achieves a competitive ratio below $3 - \frac{2}{1+|\sigma|/(4\alpha)}$.*

Proof. Consider σ to be a sequence of communication requests, divided into the following *phases*: For each phase p, the adversary chooses two arbitrary nodes u, v currently located in different

clusters. Then, the adversary generates requests between u and v, i.e., it appends them to the request sequence σ, until On moves u and v to the same cluster. At this time, the phase ends. We refer to the number of requests issued in phase p by w_p. We observe that an algorithm which never collocates the nodes cannot have a finite competitive ratio, and hence, phases and w_p are finite. After phase p, the next phase starts, and the adversary issues w_{p+1} requests between nodes $\{v, z\}$: here, z is the (unique) node which was in the same cluster as v at phase $p-1$. Denote by λ the number of phases in σ and, for simplicity of presentation, let λ be an even number.

We observe that the strategy of any online algorithm On can be characterized by its choice of w_p values and that the cost of On for this sequence is $\text{On}(\sigma) = \sum_{p=1}^{\lambda}(w_p + 2\alpha)$. To obtain a closed-form expression on the competitive ratio, one should compare the cost of On to an optimal offline algorithm. Nevertheless, since our aim is to derive a lower bound, we can compare On to arbitrary algorithms, not necessarily optimal ones. We will compare On's cost to the following three offline algorithms: (i) The first algorithm performs no moves whatsoever and hence pays cost $\sum_{p=1}^{\lambda/2} w_{2p-1}$. (ii) The second algorithm moves only in the first phase, and pays $\sum_{p=1}^{\lambda/2} w_{2p} + 2\alpha$. (iii) The third algorithm moves at the beginning of each phase, and hence pays $\sum_{p=1}^{\lambda} 2\alpha$. In order to minimize the maximal competitive ratio against any of these three offline algorithms, On must choose $\boldsymbol{w} = \{w_1, \ldots, w_T\}$ conservatively. Let $\bar{w}_{odd} = \frac{\sum_{p=1}^{\lambda/2} w_{2p-1}}{\lambda/2}$ and $\bar{w}_{even} = \frac{\sum_{p=1}^{\lambda/2} w_{2p}}{\lambda/2}$. Then we have

$$\begin{aligned}\rho(\text{On}) &= \min_{\boldsymbol{w}} \frac{\sum_{p=1}^{\lambda}(w_p + 2\alpha)}{\min\left\{\sum_{p=1}^{\lambda/2} w_{2p-1},\ \sum_{p=1}^{\lambda/2} w_{2p} + 2\alpha,\ \sum_{p=1}^{\lambda} 2\alpha\right\}} \\ &= \min_{\bar{w}_{even}, \bar{w}_{odd}} \frac{(\lambda/2)(\bar{w}_{even} + \bar{w}_{odd}) + 2\alpha\lambda}{\min\left\{(\lambda/2)\bar{w}_{even},\ (\lambda/2)\bar{w}_{odd} + 2\alpha,\ 2\alpha\lambda\right\}}.\end{aligned}$$

Dividing by $\lambda/2$, we find that

$$\rho(\text{On}) = \min_{\bar{w}_{even}, \bar{w}_{odd}} \frac{\bar{w}_{even} + \bar{w}_{odd} + 4\alpha}{\min\left\{\bar{w}_{even},\ \bar{w}_{odd} + 4\alpha/\lambda,\ 4\alpha\right\}} \geq 3 - \lambda^{-1}$$

and the minimum is achieved for $\bar{w}_{even} = 4\alpha$ and $\bar{w}_{odd} = 4\alpha\,(1 - \lambda^{-1})$. The claim follows by noticing that after λ phases, $|\sigma| = 4\alpha\,(2\lambda - 1)$ requests have been issued.

A.2 Proof of Theorem 2

Given a deterministic online algorithm On, we specify an initial and a final node partitioning, as well as a sequence of requests σ divided into phases. Set $\ell = 2$ and name nodes as $u_1, u_2, \ldots, u_k$ and $v_1, v_2, \ldots, v_k$. Initially, the adversary places the nodes as $P_1^{(start)} = \{v_1, u_2, \ldots, u_k\}$ and $P_2^{(start)} = \{u_1, v_2, \ldots, v_k\}$, respectively. In our setting, Off can perform a single swap to reach the final partitioning $P_1^{(final)} = \{u_1, u_2, \ldots, u_k\}$ and $P_2^{(final)} = \{v_1, v_2, \ldots, v_k\}$.

The basic idea is to force On to perform at least $k-1$ swaps before reaching the final partitioning. The sequence of requests σ is defined in terms of phases, indexed by $1 \leq p \leq k-1$. Phase p involves requests between $u_1, u_2, \ldots u_{p+1}$ in a round robin manner, with the first request being $\{u_p, u_{p+1}\}$, the second $\{u_{p-1}, u_{p+1}\}$, and so on. The phase ends when On performs at least one swap, and all nodes $u_1, u_2, \ldots u_{p+1}$ are in the same cluster. We observe that any online algorithm with a finite competitive ratio eventually must perform these swaps.

Let w_p denote the number or requests (i. e., On's waiting time) in phase p. It can easily be shown that any online algorithm with finite w_p has the following properties at the end of any phase $p < k$: (i) Nodes $u_1, u_2, \ldots u_{p+1}$ are in the same cluster, whereas nodes u_i with $i > p+1$ have not been involved in any communication request so far. (ii) Node u_{p+2} is not in the same cluster as $u_1, u_2, \ldots, u_{p+1}$. The first property follows by construction. The second property follows from the fact that, during phase p, On had to swap nodes that never communicated. We can therefore assume w. l. o. g. that On swapped u_{p+2} (unless it was already on the other cluster). Since all phases end with the above properties, the best strategy for On is to reach the next phase with a minimum number of swaps, i. e., using a single swap.

We have so far established that in the best case, On reaches the final partitioning after $k-1$ phases, each involving one swap. Next we consider the number of requests w_p On served remotely during each phase p. To begin with, one can directly deduce that $w_p \geq 1$: any online algorithm that swaps before receiving at least one request has an unbounded competitive ratio. To force On to serve at least 2α remote requests during the first phase, we enhance the adversary's strategy as follows:

If $w_1 < 2\alpha$ the adversary (i) changes the final partitioning to be $C_1^{(end)} = \{v_1, u_2, \ldots, u_k\}$, $C_2^{(end)} = \{u_1, v_2, \ldots, v_k\}$, and (ii) does not issue any requests involving u_1 in phases $p > 1$. On the other hand, if $w_1 \geq 2\alpha$ the adversary's strategy remains unchanged. By the same argument used previously, in each phase p, On places nodes $u_2, \ldots u_{p+1}$ in a different cluster than node u_{p+2}, and is again forced to swap $k-1$ times. On the other hand, depending on $w_1 \geq 2\alpha$, Off either performs one swap paying 2α, or does not swap and pays w_1 communication cost. The competitive ratio is therefore

$$\rho(\text{On}) \geq \frac{\sum_{p=1}^{k-1}(w_p + 2\alpha)}{\min\{w_1, 2\alpha\}} \geq \frac{2\alpha\,(k-1)}{2\alpha} + \frac{w_1}{\min\{w_1, 2\alpha\}} + \frac{\sum_{p=2}^{k-1} w_p}{2\alpha} \geq k + \frac{k-2}{2\alpha},$$

matching our claim.

A.3 Proof of Theorem 3

We prove the claim by constructing a request sequence σ which does not reveal any information about the locations and actions of an offline algorithm Off. In such a scenario, in order to ensure a finite competitive ratio, the online algorithm On is forced to iterate through a large set of possible partitionings. Off on the other hand can settle on one partitioning, paying only for a small subset of the requests.

To demonstrate the lower bound, it is sufficient to consider requests arising from a ring graph $G = (V, E)$ of n nodes and n edges $e_i = \{u_i, u_{i+1}\}$ for $i = 1, \ldots, n-1$ and $e_n = \{u_n, u_1\}$. For every possible partitioning p of G, the *cut-edge* set $\text{cut}(p)$ contains all edges of G which are cut by the partition, i. e., $e_i \in \text{cut}(p)$ if u_i and u_{i+1} are not in the same cluster.

To define the request sequence σ, suppose that at time t, On has chosen a partitioning p_t and let $\text{cut}(p_t)$ be the corresponding cut-edge set. The adversary then issues the next request across the ring cut-edge adjacent to the node with the smallest id, formally $\sigma_{t+1} = \arg\min_{e_i \in \text{cut}(p_t)} i$. By construction, σ has two important properties: (i) On *will have to serve every request remotely.* (ii) σ *is uniquely determined by the online algorithm and does not depend on* Off*'s actions.* Accordingly, the evolution of the request sequence is entirely under the control of On.

By construction, On's cost in each time step is at least 1 (serve remotely or swap), so its cost up to time t is $\text{On}(\sigma_{1t}) \geq t$, where $\sigma_{1t} = \{\sigma_1, \ldots, \sigma_t\}$. Regarding Off's cost, let an *order-preserving* partitioning be a partitioning in which nodes in the same cluster have contiguous ids on the ring graph. Since there are k possible order-preserving partitionings o_m, where $m = 1, \ldots, k$, we define for each partitioning the corresponding cut-edge sets as $\text{cut}(o_m) = \{e_i : (i-1) \mod k = m-1\}$. It is crucial to observe that different order-preserving cut-edge sets are disjoint, meaning that $\text{cut}(o_m) \cap \text{cut}(o_{m'}) = \varnothing$ for all $m \neq m'$. Before the first request is issued, Off chooses an order-preserving partitioning o and performs the necessary swaps to reach it, paying migration costs of $\text{mig}(\sigma_1; \text{Off}) \leq \alpha \cdot n$. Off does not swap nodes again. The key point here is that, as far as On is concerned, Off might have chosen any order-preserving partitioning. We now claim that Off's communication cost is $\text{Off}(\sigma_{1t}) \leq t/k$. Based on On's actions, Off chooses to migrate to an order-preserving partitioning o^* that minimizes its communication cost. Let $\text{com}(\sigma_{1t}, o_m; \text{Off})$ denote the *communication* cost of Off if it had chosen to migrate to partitioning o_m. Since the order-preserving cuts are disjoint, we have that at time t, $\sum_{m=1}^{k} \text{com}(\sigma_{1t}, o_m; \text{Off}) = t$. Therefore, clearly $\min_m \text{com}(\sigma_{1t}, o_m; \text{Off}) \leq t/k$ implying also that $\text{com}(\sigma_{1t}; \text{Off}) \leq t/k$.

It easily follows that, as the length of the request sequence grows, the competitive ratio approaches k.

$$\begin{aligned} \rho(\text{On}) = \lim_{t\to\infty} \frac{\text{On}(\sigma_{1t})}{\text{Off}(\sigma_{1t})} &\geq \lim_{t\to\infty} \frac{t}{\text{com}(\sigma_{1t}; \text{Off}) + \text{mig}(\sigma_{1t}; \text{Off})} \\ &\geq \lim_{t\to\infty} \frac{t}{t/k + \alpha \cdot n} = k \end{aligned} \tag{3}$$

To see that the above holds independently of the augmentation (given that $\delta < \ell$) notice that, as long as On cannot fit the entire ring within a single cluster ($k' < k \cdot \ell$), the cut edge $\text{cut}(p_t)$ sets will be non-empty. The adversary can therefore always generate the request sequence as above, forcing On to serve each request remotely, and the argument proceeds as before.

A.4 Proof of Lemma 3

Consider for now a scenario where there exists only a *last epoch*, and let us hence focus on the case where CREP and Off start from the same initial configuration. We analyze the cost induced by

the communication graph G, at some time t. We can bound the cost of Off as follows: Let Off place the components of the graph in a statically optimal way, to minimize the communication cost between clusters (i.e., a graph G partition).

Clusters. We start by examining the costs pertaining to clusters, while ignoring edges between components (for now). There could be two types of clusters with regard to the initial configuration: *identical* and *unidentical* clusters. An identical cluster s is a cluster which has the same nodes as the initial configuration.

Identical clusters. For any identical cluster s there are two options, either CREP *touched* cluster s, i.e., CREP migrates at least one node in or out from s, or not. If the cluster was not touched by CREP, then Off and CREP pay the same cost for this cluster. If CREP touched the cluster, then Off's communication cost is at least α for s. We now bound the cost of CREP for s: recall from Property 2 that for each component ϕ, the cost of migration and internal communication of CREP for ϕ is at most $2\alpha \cdot (|\phi| \log |\phi|)$. Hence, the cost for migration and internal communication of all components in a cluster is bounded by $2\alpha \cdot (k \log k)$.

Unidentical clusters. Consider a cluster which is *unidentical* to its initial configuration. The cost incurred by Off on this cluster is at least 2α, since it has to perform at least one swap to arrange it form the initial configuration. On the other hand, the total cost of CREP, over all components in the cluster, is again bounded by $2\alpha \cdot (k \log k)$.

Crossing edges. It remains to bound the cost of the edges between components in G. We know that while CREP paid the full price $\sum_{ij} w_{ij}$ for these edges, Off may have been able to arrange many edges within the same cluster and avoid their cost. We show next that this cannot happen. We call an edge between two components (ϕ_i, ϕ_j) *light* if $|\phi_i| + |\phi_j| \leq k$, and correspondingly *heavy* if $|\phi_i| + |\phi_j| > k$.

Light edges. Consider a set of light edges that Off put in the same cluster s. There could be two types of components connected to these edges: *broken* and *unbroken* ones. A broken component is a component which Off did not place entirely within s, but split between multiple clusters. Let B denote the set of broken components. Correspondingly, an unbroken component is a component which is completely within s, and U is the set of unbroken components. The cost of light edges between unbroken components can be at most $|U| - 1 \leq k - 1$, times 2α since, otherwise, U would have been merged by CREP to a new component. For each broken component, Property 3 entails a communication cost of at least α to nodes outside s, which Off must pay. Each of these components can have at most $|U|$ light edges to components in U, entailing a total cost of at most $k\alpha$.

Heavy edges. If Off puts a heavy edge with $k < |\phi_i| + |\phi_j| \leq 2k$ into a cluster, it must split either ϕ_i or ϕ_j among two clusters, which, again by Property 3, incurs communication costs of at least α, between the two parts of the divided component. Since this holds for any partition, any swap made by Off will only increase the lower bound.

Concluding the first part of our proof, the competitive ratio holds as, for every cluster and its components, the ratio is $O(k \log k)$. For light edges the ratio is $O(k)$ and for heavy edges the ratio is $O(k)$.

Multiple epochs. Finally, we relax the assumption that there is only a *last epoch*. So far, we assumed that both Off and CREP start the last epoch from the same initial configuration. We find that, differently from the above, Off now does not necessarily pay for the components of touched clusters; it can exploit the migrations induced by previous epochs to rearrange components in an optimal way. Nevertheless, for every difference in the configuration of CREP and Off in the beginning of the last epoch, there must exist a point in the past at which Off paid at least α (either because Off performed a migration, or because CREP migrated whereas Off did not). By charging Off half the cost in previous epochs, we can charge Off at least cost $\alpha/2$ for every touched cluster during the last epoch. The competitive ratio in Lemma 3 thus also holds here.

A.5 Simulation of Online Paging

For $\ell = 2$, BRP without augmentation is a generalization of online paging, and hence the linear-in-k lower bounds for online paging also apply to BRP. To see this, consider a variant of online paging with k items, cache size $k - 1$, and input sequence σ. We can simulate this problem in BRP, with request sequence σ', as follows: the first cluster (size k) will be used to represent the cache, while the second cluster (also size k) will be used to represent the slow memory. The $k - 1$ first items in the first cluster will be the cache content, and in addition, the cluster contains a dummy item

d; the other items are irrelevant. Whenever an item i is requested in σ, we insert $2\alpha + 1$ requests to σ' between i and d. In order to "cache" i, i.e., move i to the first cluster, a deterministic online algorithm must "evict" some other item, i.e., move it to the second cluster. In order to ensure that the dummy item d is never moved to the second cluster and hence keeps blocking the slot for the k-th item, for any two consecutive requests in σ, α additional requests are introduced in σ' between any current item in the first cluster and d. Thus, the eviction problem can be simulated also in BRP.